

Docket No. PD-200053
CUSTOMER NO.: 020991

Please charge Deposit Account No. 50-0383 of Hughes Electronics Corporation, El Segundo, California, in the amount of **\$1,410.00**. The Commissioner is hereby authorized to charge any additional fees which may be required, or credit any overpayment, to that account.

The Commissioner is further hereby authorized to charge to said above Deposit Account No. 50-0383, pursuant to 37 CFR 1.25(b), any fees whatsoever which may properly become due or payable, as set forth in 37 CFR 1.16 to 1.17 inclusive, for the entire pendency of this application without specific additional authorization.

Please associate this application with the Hughes Electronics Corporation Customer Number 020991.

This form is submitted in triplicate.

HUGHES ELECTRONICS CORPORATION



John T. Whelan
Reg. No.: 32,448

HUGHES ELECTRONICS CORPORATION
Patent Docket Administration
Bldg. 001, M/S A109
PO Box 956
El Segundo, CA 90245-0956
Telephone: (301) 601-7252
Date: September 15, 2000

PERFORMANCE ENHANCING PROXY AND METHOD FOR
ENHANCING PERFORMANCE

CROSS-REFERENCE TO RELATED APPLICATIONS

This application claims the benefit under 35 U.S.C. § 119(e) of U.S. Provisional Application of John Border et al. entitled "Enhanced TCP Spoofing", serial no. 60/185,553, filed on February 28, 2000, and a U.S. Provisional Application of John Border et al. entitled "Performance Enhancing Proxy", serial no. 60/220,026, filed on July 21, 2000, the entire contents of both are incorporated by reference herein.

The present application is also related to co-pending application in the name of Matt Butehorn et al., entitled "Selective Spoofer and Method of Performing Selective Spoofing", filed concurrently herewith.

BACKGROUND OF THE INVENTION

The present invention is generally directed to a method and apparatus for improving the performance of protocols on network paths, and more particularly, a method and apparatus for improving the performance of the TCP/IP protocol on the Internet, utilizing a performance enhancing proxy.

DESCRIPTION OF THE RELATED ART

The transmission control protocol (TCP) is the dominant protocol in use today on the Internet. TCP is carried by the Internet protocol (IP) and is used in a variety of applications including reliable file transfer and Internet web page access applications. The four layers of the TCP/IP protocol suite are illustrated in Fig. 1. As illustrated, the link layer (or the network interface layer) includes device drivers in the operating system and any corresponding network interface cards. Together, the device driver and

the interface cards handle hardware details of physically interfacing with any cable or whatever type of media is being used. The network layer (also called the Internet layer) 12 handles the movement of packets around the network. Routing of packets, for example, takes place at the network layer 12. IP, Internet control message protocol (ICMP), and Internet group management protocol (IGMP) may provide the network layer in the TCP/IP protocol suite. The transport layer 14 provides a flow of data between two hosts, for the application layer 16 above.

In the TCP/IP protocol suite, there are at least two different transport protocols, TCP and a user datagram protocol (UDP). TCP, which provides a reliable flow of data between two hosts, is primarily concerned with dividing the data passed to it from the application layer 16 into appropriately sized chunks for the network layer 12 below, acknowledging received packets, setting timeouts to make certain the other end acknowledges packets that are sent, and so on. Because this reliable flow of data is provided by the transport layer 14, the application layer 16 can ignore these details. UDP, on the other hand, provides a much simpler service to the application layer 16. UDP just sends packets of data called datagrams from one host to another, but there is no guarantee that the datagrams reach the other end. Any desired reliability must be added by the application layer 16.

The application layer 16 handles the details of the particular application. There are many common TCP/IP applications that almost every implementation provides. These include telnet for remote log-in, the file transfer protocol (FTP), the simple mail transfer protocol (SMTP) or

electronic mail, the simple network management protocol (SNMP), and many others.

As described above, TCP provides reliable, in-sequence delivery of data between two IP hosts. The IP hosts set up a TCP connection, using a conventional TCP three-way handshake and then transfer data using a window based protocol with the successfully received data acknowledged. Fig. 2 illustrates an example of the conventional TCP three-way handshake between IP hosts 20 and 22. First, the IP host 20 that wishes to initiate a transfer with IP host 22, sends a synchronize (SYN) signal to IP host 22. The IP host 22 acknowledges the SYN signal from IP host 20 by sending a SYN acknowledgement (ACK). The third step of the conventional TCP three-way handshake is the issuance of an ACK signal from the IP host 20 to the IP host 22. IP host 22 is now ready to receive the data from IP host 20 (and vice versa). After all the data has been delivered, another handshake (similar to the described to initiate the connection) is used to close the TCP connection.

TCP was designed to be very flexible and works over a wide variety of communication links, including both slow and fast links, high latency links, and links with low and high error rates. However, while TCP (and other high layer protocols) works with many different kinds of links, TCP performance, in particular, the throughput possible across the TCP connection, is affected by the characteristics of the link in which it is used. There are many link layer design considerations that should be taken into account when designing a link layer service that is intended to support Internet protocols. However, not all characteristics can be compensated for by choices in the link layer design. TCP

has been designed to be very flexible with respect to the links which it traverses.

An alternative to a tailored protocol is the use of performance enhancing proxies (PEPs), to perform a general class of functions termed "TCP spoofing," in order to improve TCP performance over impaired (i.e., high latency or high error rate) links. TCP spoofing involves an intermediate network device (the performance enhancing proxy (PEP)) intercepting and altering, through the addition and/or deletion of TCP segments, the behavior of the TCP connection in an attempt to improve its performance.

Conventional TCP spoofing implementations include the local acknowledgement of TCP data segments in order to get the TCP data sender to send additional data sooner than it would have sent if spoofing were not being performed, thus improving the throughput of the TCP connection. Generally, conventional TCP spoofing implementations have focused simply on increasing the throughput of TCP connections either by using larger windows over the link or by using compression to reduce the amount of data which needs to be sent, or both.

Many TCP PEP implementations are based on TCP ACK manipulation. These may include TCP ACK spacing where ACKs which are bunched together are spaced apart, local TCP ACKs, local TCP retransmissions, TCP ACK filtering and reconstruction. Other PEP mechanisms include tunneling, compression, and priority-based multiplexing.

SUMMARY OF THE INVENTION

The present invention is directed to a method and apparatus for enhancing the performance of a network.

The performance enhancing functions of the present invention are applicable to a wide variety of communication links, including both slow and fast links, high latency links, and links with low and high error rates.

The performance enhancing functions, which may be implemented either singly or in combination include:

- selective TCP spoofing which allows flexible configuration of which connections should be spoofed;
- spoofing of the conventional TCP three-way handshake;
- local data acknowledgement, which allows data windows to increase at local speeds;
- multiplexing multiple connections across a single connection which increases acknowledgement traffic reduction and provides a backbone connection protocol optimized for the particular backbone link being used;
- data compression/encryption;
- prioritizing access to the link; and
- selecting a particular path for the data accompanying a connection to take.

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 illustrates the four layers of the conventional TCP/IP protocol suite.

Figure 2 illustrates an example of the conventional TCP three-way handshake between IP hosts.

Figure 3 illustrates an exemplary network in which the performance enhancing proxy (PEP) of the present invention is implemented.

Fig. 4 illustrates an exemplary embodiment of a performance enhancing proxy (PEP), in one exemplary embodiment of the present invention.

Fig. 5 illustrates an exemplary stack, which describes the relationship between the conventional TCP stack and the PEP kernels, in one exemplary embodiment of the present invention.

Figure 6 illustrates the PEP kernels in one exemplary embodiment of the present invention

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

Fig. 3 illustrates an exemplary network 100 in which the performance enhancing proxy (PEP) of the present invention may be utilized. The network 100 in Fig. 3 includes a plurality of hosts 110 connected to a network gateway 120 via TCP connections. The network gateway 120 is connected to another network gateway 140 via a backbone connection on a backbone link 130. In Fig. 3, the backbone link 130 is depicted as a satellite link, however this is only exemplary. The network gateway 140 is further connected to a second group of hosts 150, also via TCP connections. In the arrangement illustrated in Fig. 3, the network gateways 120, 140 facilitate communication between the groups of hosts 110, 150. The network gateways 120, 140 facilitate communication between the two groups of hosts 110, 150 by performing the following performance enhancing functions, either singly or in combination;

- selective TCP spoofing which allows flexible configuration of which TCP connections should be spoofed;
- spoofing of the conventional TCP three-way handshake, with the TCP connections terminated at each end of the backbone link 130;
- local data acknowledgement, allowing TCP windows to increase at local speeds;

- multiplexing multiple TCP connections across a single backbone connection which provides the following benefits,
- increased acknowledgement traffic reduction with data from multiple TCP connections acknowledged by a single backbone connection acknowledgement,
- support for high throughput TCP connections using a backbone connection protocol optimized for the particular backbone link being used;
- compression of data sent over the backbone link 130 to reduce the amount of traffic to be sent, further leveraging the capabilities of the backbone connection;
- encryption of data sent over the backbone link 130 to protect data privacy;
- prioritized access to backbone link 130 capacity on a per TCP connection basis; and
- selecting a particular path for the data accompanying a connection to take.

These features are described in more detail below.

Fig. 4 illustrates an exemplary embodiment of a performance enhancing proxy 200, the functionality of which can be implemented in a network gateway 120, 140 in one exemplary embodiment. The PEP 200 includes a platform environment 210, which includes the hardware and software operating system. The PEP 200 also includes local area network (LAN) interfaces 220 and wide area network (WAN) interfaces 230. In the example in Fig. 3, the network gateway 120 may establish the TCP connections with the IP hosts 110, via a local LAN interface 220 and may establish the backbone connection with the network gateway 140 via a WAN interface 230. The PEP platform environment 210 may

also include general elements which perform the functions of routing 240, buffer management 250, event management 260, and parameter management 270. As illustrated in Fig. 4, the network gateway also includes a TCP spoofing kernel (TSK) 280, a backbone protocol kernel (BPK) 282, a prioritization kernel (PK) 284, and a path selection kernel (PSK) 286. These four kernels essentially make up the functionality of the performance enhancing proxy 200.

The platform environment 210 has at least three purposes. These include performing functions that the various PEP kernels 280, 282, 284, 286 cannot perform directly because the implementation of the function is platform specific. This arrangement has the advantageous effect of hiding platform specific details from the PEP kernels 280, 282, 284, 286. An example of a platform specific function is the allocation of a buffer. In some platforms, buffers are allocated, while in other platforms, buffers are created at start-up and organized into linked lists. It is noted that platform specific functions are not limited to functions generic to all of the kernels 280, 282, 284, 286.

The second purpose of the platform environment 210 is to provide the task context in which the PEP kernels 280, 282, 284, 286 run. In one exemplary embodiment, all PEP kernels 280, 282, 284, 286 can run in the same task context for efficiency, but this is not required.

The third purpose of the platform environment 210 is to provide an interface between the PEP functionality (embodied in kernels 280, 282, 284, 286) and the other functionality of the network gateway 120, 140. For example, the platform environment 210 may provide the interface between the PEP functionality and the routing function 240 illustrated in

Fig. 4. It is noted that the platform specific functions illustrated in Fig. 4 are examples and are not considered an exhaustive list. It is further noted that the PEP kernels shown touching each other (280, 282 and 284, 286) in Fig. 4 may have a direct procedural interface to each other. Further, the kernels 280, 282, 284, 286 may include direct interfaces, as opposed to routing everything through the platform environment 210 (as shown in Figure 4), in order to improve performance.

Fig. 5 is an exemplary stack, which illustrates the relationship between the TCP stack and the PEP kernels 280, 282, 284, 286 of the present invention. The TSK 280 is primarily responsible for functions related to TCP spoofing. The TSK 280 includes at least two basic elements, the conventional transport layer 14 of the TCP/IP stack (such as the transport layer 14 illustrated in Fig. 1) and a TCP spoofing application 2802. The transport layer 14 is responsible for interacting with the TCP stacks of IP hosts 110 connected to a PEP's local LAN interface 220. The TSK 280 implements the TCP protocol including the appropriate TCP state machines and terminates spoofed TCP connections. The TCP spoofing application 2802 can sit on top of the transport layer 14 and act as the application receiving data from and sending data to the IP hosts 110 applications. The TCP spoofing application 2802 can hide the details of TCP spoofing from the transport layer 14, allowing the transport layer 14 to function as much like a standard transport layer 14 as possible. The TCP spoofing application 2802 can also interface to the BPK 282.

The PK 284 is responsible for determining the priority of IP packets and then allocating transmission opportunities based on priority. The PK 284 can also control access to

buffer space by controlling the queue sizes associated with sending and receiving IP packets.

The PSK 286 determines which path an IP packet should take to reach its destination. The path selected by the PSK 286 can be determined applying path selection rules. PSK 286 may also determine which IP packet should be forwarded using an alternate path and which packets should be dropped when one or more primary paths fail.

The BPK 282 performs backbone protocol maintenance and is the protocol by which the network gateways 120, 140 in Fig. 3 communicate. The BPK 282 provides reliable delivery of data, uses a relatively small amount of acknowledgement traffic, and supports generic backbone use (i.e., use not specific to the TSK 280). One such example is the reliable data protocol (RDP).

Figure 6 illustrates the PEP kernels 280, 282, 284, and 286 as well as other kernels such as a data compression kernel 290 and an encryption kernel 292. As described above, the PEP kernels 280, 282, 284, and 286 facilitate communication between the two groups of hosts 110, 150, by performing a variety of performance enhancing functions, either singly or in combination. These performance enhancing functions are described in more detail below.

Selective TCP Spoofing

Selective TCP Spoofing is performed by the TSK 280 and includes a set of user configurable rules that are used to determine which TCP connections should be spoofed. Selective TCP spoofing improves performance by not tying up TCP spoofing-related resources, such as buffer space, control blocks, etc., for TCP connections for which the user has determined that spoofing is not beneficial or required.

In particularly, the TSK 280 discriminates among the various TCP connections based on the applications using them. TCP spoofing is then performed only for those TCP connections that are associated with applications for which high throughput or reduced connection startup latency (or both) is required. As a result, the TSK 280 conserves TCP spoofing resources for only those TCP connections for which high throughput or reduced connection startup latency (or both) is required. Further, the TSK 280 increases the total number of TCP connections which can be active before running out of TCP spoofing resources, since any active TCP connections which do not require high throughput are not allocated resources.

One criterion for identifying TCP connections of applications for which TCP spoofing should and should not be performed is the TCP port number field contained in the TCP packets being sent. In general, unique port numbers are assigned to each type of application. Which TCP port numbers should and should not be spoofed can be stored in the TSK 280. The TSK 280 is also re-configurable to allow a user or operator to reconfigure the TCP port numbers which should and should not be spoofed. The TSK 280 also permits a user or operator to control which TCP connections are to be spoofed based on other criteria. In general, a decision on whether to spoof a TCP connection may be based on any field within a TCP packet. The TSK 280 permits a user to specify which fields to examine and which values in these fields identify TCP connections that should or should not be spoofed. Another example of a potential use for this capability for the user or operator to select the IP address of the TCP packet in order to control for which users TCP spoofing is performed. The TSK 280 also permits a user to

look at multiple fields at the same time. As a result, the TSK 280 permits a user or operator to use multiple criteria for selecting TCP connections to spoof. For example, by selecting both the IP address and the TCP port number
5 fields, the system operator can enable TCP spoofing for only specific applications from specific users.

The user configurable rules may include five exemplary criteria which can be specified by the user or operator in producing a selective TCP spoofing rule:

- 10 • Destination IP address;
- Source IP address;
- TCP port numbers (which may apply to both the TCP destination and source port numbers);
- TCP options; and
- 15 • IP differentiated services (DS) field.

As outlined above, in addition to supporting selective TCP spoofing rules for each of these criterion, AND and OR combination operators can be used to link criteria together. For example, using the AND combination operator, a rule can
20 be defined to disable TCP spoofing for FTP data received from a specific host. Also, the order in which the rules are specified may be significant. It is possible for a connection to match the criteria of multiple rules. Therefore, the TSK 280 can apply rules in the order
25 specified by the operator, taking the action of the first rule that matches. A default rule may also be set which defines the action to be taken for TCP connections which do not match any of the defined rules. The set of rules selected by the operator may be defined in a selective TCP
30 spoofing selection profile.

As an example, assume enough buffer space has been allocated to spoof five (5) TCP connections. If four (4)

low speed applications (i.e. applications which, by their nature, do not require high speed) bring up connections along with one high speed application, the high speed connection has access to only 1/5 of the available spoofing buffer space. Further, if five (5) low speed connections are brought up before the high speed connection, the high speed connection cannot be spoofed at all. Using the TSK 280 selective spoofing mechanism, the low speed connections are not allocated any spoofing buffer space. Therefore, the high speed connection always has access to all of the buffer space, improving its performance with respect to an implementation without the selective TCP spoofing feature of the TSK 280.

Three-Way Handshake Spoofing

The TSK 280 also facilitates spoofing of the conventional three-way handshake. Three-Way Handshake Spoofing involves locally responding to a connection request to bring up a TCP connection in parallel with forwarding the connection requests across the backbone link 130. This allows the originating IP host (for example, 110) to reach the point of being able to send the data it must send at local speeds, i.e. speeds that are independent of the latency of the backbone link 130. Three-Way Handshake Spoofing allows the data that the IP host 110 needs to send to be sent to the destination IP host 150 without waiting for the end-to-end establishment of the TCP connection. For backbone links 130 with high latency, this significantly reduces the time it takes to bring up the TCP connection and, more importantly, the overall time it takes to get a response (from an IP host 150) to the data the IP host 110 sends.

A specific example where this technique is useful relates to an Internet web page access application. With three-way handshake spoofing, an IP host's request to retrieve a web page can be on its way to a web server without waiting for the end-to-end establishment of the TCP connection, thereby reducing the time it takes to download the web page.

Local Data Acknowledgement

With Local Data Acknowledgement, the TSK in the network gateway 120 (for example) locally acknowledges data segments received from the IP host 110. This allows the sending IP host 110 to send additional data immediately. More importantly, TCP uses received acknowledgements as signals for increasing the current TCP window size. As a result, local sending of the acknowledgements allows the sending IP host 110 to increase its TCP window at a much faster rate than supported by end to end TCP acknowledgements. The TSK 280 (the spoofer) takes on the responsibility for reliable delivery of the data which it has acknowledged.

TCP Connection to Backbone Connection Multiplexing

In the BPK 282, multiple TCP connections are multiplexed onto and carried by a single backbone connection. This improves system performance by allowing the data for multiple TCP connections to be acknowledged by a single backbone connection acknowledgement (ACK), significantly reducing the amount of acknowledgement traffic required to maintain high throughput across the backbone link 130. In addition, the BPK 282 selects a backbone connection protocol that is optimized to provide high throughput for the particular link. Different backbone connection protocols can be used by the BPK 282 with

different backbone links without changing the fundamental TCP spoofing implementation. The backbone connection protocol selected by the BPK 282 provides appropriate support for reliable, high speed delivery of data over the backbone link 130, hiding the details of the impairments (for example high latency) of the link from the TCP spoofing implementation.

The multiplexing by the BPK 282 allows for the use of a backbone link protocol which is individually tailored for use with the particular link and provides a technique to leverage the performance of the backbone link protocol with much less dependency upon the individual performance of the TCP connections being spoofed than conventional methods.

Further, the ability to tailor the backbone protocol for different backbone links makes the present invention applicable to many different systems. These include multimedia networks such as the DirecWay™ Multimedia Network, the Integrated Satellite Business Networks™ (ISBN™), other types of VSAT Networks, and TDMA Networks.

Data Compression/Encryption

The PEP 200 may also include a data compression kernel 290 for compressing TCP data and an encryption kernel 292 for encrypting TCP data. Data compression increases the amount of data that can be carried across the backbone connection. Different compression algorithms can be supported by the data compression kernel 290 and more than one type of compression can be supported at the same time. The data compression kernel 290 may optionally apply compression on a per TCP connection basis, before the TCP data of multiple TCP connections is multiplexed onto the backbone connection or on a per backbone connection basis,

after the TCP data of multiple TCP connections has been multiplexed onto the backbone connection. Which option is used is dynamically determined based on user configured rules and the specific compression algorithms being utilized. Exemplary data compression algorithms are disclosed in U.S. Patent Nos. 5,973,630, 5,955,976, the entire contents of which are hereby incorporated by reference. The encryption kernel 292 encrypts the TCP data for secure transmission across the backbone link 130.

Encryption may be performed by any conventional technique. It is also understood that the corresponding spoofer (in the example outlined above, the network gateway 140) includes appropriate kernels for decompression and decryption, both of which may be performed by any conventional technique.

Prioritization

The PK 284 provides prioritized access to the backbone link 130 capacity. For example, the backbone connection can actually be divided into N ($N > 1$) different sub-connections, each having a different priority level. In one exemplary embodiment, four priority levels can be supported. The PK 284 uses user-defined rules to assign different priorities, and therefore different sub-connections of the backbone connection, to different TCP connections. The PK 284 also uses user-defined rules to control how much of the backbone link 130 capacity is available to each priority level. There are at least six criteria which can be used to determine priority:

- Destination IP address;
- Source IP address;
- IP next protocol;

- TCP port numbers (which may apply to both the TCP destination and source port numbers);
- UDP port numbers (which may apply to both the UDP destination and source port numbers); and
- 5 • IP differentiated services (DS) field.

The type of data in the TCP data packets may also be used as a criterion. For example, video data could be given highest priority. Mission critical data could also be given high priority.

10 As outlined above, in addition to supporting selective prioritization rules for each of these criteria, AND and OR combination operators can be used to link criteria together. For example, using the AND combination operator, a rule can be defined to assign a priority for FTP data received from a
15 specific host. Also, the order in which the rules are specified may be significant. It is possible for a connection to match the criteria of multiple rules. Therefore, the PK 284 can apply rules in the order specified by the operator, taking the action of the first rule that
20 matches. A default rule may also be set which defines the action to be taken for TCP connections which do not match any of the defined rules. The set of rules selected by the operator may be defined in a prioritization profile.

Path Selection

25 The PSK 286 is responsible for determining which path an IP packet should take to reach its destination. The path selected by the PSK 286 can be determined by applying path selection rules. The PSK 286 also determines which IP packets should be forwarded using an alternate path and
30 which IP packets should be dropped when one or more primary paths fail. Path selection parameters can also be

configured using profiles. The path selection rules may be designed to provide flexibility with respect to assigning paths while making sure that all of the packets related to the same traffic flow (e.g., the same TCP connection) take the same path (although it is also possible to send segments of the same TCP connection via different paths, this segment "splitting" may have negative side effects). There are at least seven criteria which can be used to select a path:

- priority of the IP packet as set by the PK 284 (should be the most common criterion)
- Destination IP address;
- Source IP address;
- IP next protocol;
- TCP port numbers (which may apply to both the TCP destination and source port numbers);
- UDP port numbers (which may apply to both the UDP destination and source port numbers); and
- IP differentiated services (DS) field.

As outlined above, in addition to supporting path selection rules for each of these criteria, AND and OR combination operators can be used to link criteria together. For example, using the AND combination operator, a rule can be defined to select a path for FTP data received from a specific host. Also, the order in which the rules are specified may be significant. It is possible for a connection to match the criteria of multiple rules. Therefore, the PSK 286 can apply rules in the order specified by the operator, taking the action of the first rule that matches. A default rule may also be set which defines the action to be taken for TCP connections which do not match any of the defined rules. The set of rules

selected by the operator may be defined in a path selection profile.

A path selection rule may select the following path information:

- the primary path for any IP packets which match the rule. A primary path should be specified in any path selection rule (including the default rule, discussed below);
- the secondary path for any IP packets which match the rule. The secondary path should only be used when the primary path has failed. If no secondary path is specified, any IP packets which match the rule can be discarded when the primary path fails; and
- the tertiary path for any IP packets which match the rule. A tertiary path should only be specified if a secondary path is specified. The tertiary path should only be used when both the primary and secondary paths have failed. If no tertiary path is specified, any IP packets which match the rule can be discarded when both the primary and secondary paths fail.

Path selection may be generalized such that the path selection rule can select up to N paths where the Nth path is used only if the (N-1)th path fails. The example above where N=3 is merely illustrative, although N is typically a fairly small number.

The operation of the entire network is described below in conjunction with Figures 3 and 6. First, a backbone connection is established between the PEPs 200 of two network gateways 120, 140 (the two spoofers), located at each end of the backbone link 130 for which TCP spoofing is desired. Whenever an IP host 110 initiates a TCP connection, the TSK 280 of the PEP 200 local to the IP host

110 checks its configured selective TCP spoofing rules. If the rules indicate that the connection should not be spoofed, the PEP 200 allows the TCP connection to flow end-to-end unspoofed. If the rules indicate that the connection should be spoofed, the spoofing PEP 200 locally responds to the IP host's TCP three-way handshake. In parallel, the spoofing PEP 200 sends a message across the backbone link 130 to its partner network gateway 140 asking it to initiate a TCP three-way handshake with the IP host 150 on its side of the backbone link 130. Data is then exchanged between the IP host 110, 150 with the PEP 200 of the network gateway 120 locally acknowledging the received data and forwarding it across the backbone link 130 via the high speed backbone connection, compressing the data as appropriate based on the configured compression rules. The priority of the TCP connection is determined when the connection is established. The BPK 282 can multiplex the connection with other received connections over a single backbone connection, the PK 284 determines the priority of the connection and the PSK determines the path the connection is to take.

In summary, the PEP 200 described above improves network performance by allocating TCP spoofing-related resources, such as buffer space, control blocks, etc., only to TCP connections for which spoofing is beneficial; by spoofing the three-way handshake to decrease data response time; by reducing the number of ACKs which are transmitted by performing local acknowledgement and by acknowledging multiple TCP connections with a single ACK; by performing data compression to increase the amount of data that can be transmitted; by assigning priorities to different connections; and by defining multiple paths for connections to be made.

Although the present invention has been described in conjunction with exemplary Figures 3-6, the present invention may be varied in many ways obvious to one of ordinary skill in the art. For instance, although the present invention describes spoofing some connections based on their associated applications, spoofs acknowledgements (ACKs), spoofing a three-way handshake, multiplexing multiple connections onto a single connection, prioritizing connections, performing path selection, compressing, encrypting data, any other performance enhancing function known to one of ordinary skill in the art could also be implemented.

Similarly, although the various parameters discussed above in conjunction with Figures 3-6 include destination address, source address, destination port number, source port number, options, a differentiated services(DS)field, protocol, a differentiated services(DS)field, a type of data contained therein, and priority of the connection or combinations thereof, any other parameters known to one of ordinary skill in the art could also be utilized.

Further, although the present invention has been described above using the TCP, TCP/IP, or UDP protocols, any high layer protocol known to one of ordinary skill in the art could also be utilized. Although the present invention has been described above in conjunction with a satellite link, any impaired link, that is any link with at least one potentially negative parameter (high latency, high bit error rate, etc.) could also benefit from the various performance enhancing features of the present invention. Although the various performance enhancing features of the present invention have been described as taking place within a network gateway, these functions could be performed within

any network element, including, but not limited to, a host, a hub, a remote, and a router. Further, although the functionality described above in conjunction with the present invention has been described as being originally resident within a network element, the functionality may be added to an existing network element, via software loaded from an article of manufacture or software downloaded via a propagated signal.

5

What is claimed:

- 1 1. A network apparatus, comprising:
2 a performance enhancing proxy which facilitates
3 communication between said network apparatus and other
4 network entities by performing at least one performance
5 enhancing function.
- 1 2. The network apparatus of claim 1, wherein said
2 network apparatus is connected to other network entities via
3 a first type of connection and other network entities via a
4 second type of connection.
- 1 3. The network apparatus of claim 2, wherein said
2 performance enhancing proxy establishes multiple connections
3 of the first type associated with different applications,
4 said performance enhancing proxy including,
5 a spoofing element, which spoofs some of the multiple
6 connections of the first type based on their associated
7 applications.
- 1 4. The network apparatus of claim 3, wherein said
2 spoofing element only spoofs connections of the first type
3 associated with at least one of applications with high
4 throughput and applications for which reduced startup
5 latency is desired.
- 1 5. The network apparatus of claim 3, wherein said
2 spoofing element assigns spoofing resources, including
3 buffer space and control blocks, to the spoofed connections.

6. The network apparatus of claim 3, wherein said spoofing element spoofs connections using at least one spoofing rule based on destination address, source address, destination port number, source port number, options, a differentiated services (DS) field or combinations thereof.

7. The network apparatus of claim 6, wherein said spoofing element defines the at least one spoofing rule in a spoofing profile.

8. The network apparatus of claim 2, wherein said performance enhancing proxy establishes multiple connections of the first type, said performance enhancing proxy including, a spoofing element, which spoofs acknowledgements (ACKs).

9. The network apparatus of claim 2, wherein said performance enhancing proxy establishes multiple connections of the first type, said performance enhancing proxy including, a spoofing element, which spoofs a three-way handshake between said network apparatus and another network entity.

10. The network apparatus of claim 2, wherein said performance enhancing proxy establishes multiple connections of the first type, said performance enhancing proxy including, a protocol element, which multiplexes multiple connections of the first type onto a single connection of the second type.

1 11. The network apparatus of claim 2, wherein said
2 performance enhancing proxy establishes multiple connections
3 of the first type, said performance enhancing proxy
4 including,

5 a prioritization element, which prioritizes connections
6 of the first type to determine what priority level of the
7 connection of the second type, each of the connections of
8 the first type are assigned.

1 12. The network apparatus of claim 11, wherein said
2 prioritizing element prioritizes connections using at least
3 one prioritizing rule based on destination address, source
4 address, destination port number, source port number,
5 protocol, a differentiated services (DS) field, a type of
6 data contained within the connection or combinations
7 thereof.

1 13. The network apparatus of claim 12, wherein said
2 prioritizing element defines the at least one prioritizing
3 rule in a prioritizing profile.

1 14. The network apparatus of claim 2, wherein said
2 performance enhancing proxy establishes multiple connections
3 of the first type, said performance enhancing proxy
4 including,

5 a path selection element, which selects a path for data
6 associated with connections of the first type across
7 connections of the second type or connections of other
8 types.

1 15. The network apparatus of claim 14, wherein said
2 path selection element can select up to N paths ($N > 1$), where
3 the Nth path is selected only if the (N-1)th path fails.

1 16. The network apparatus of claim 15, wherein said
2 path selection element selects a path using at least one
3 path selection rule based on priority, a destination
4 address, source address, destination port number, source
5 port number, protocol, a differentiated services (DS) field
6 or combinations thereof.

1 17. The network apparatus of claim 16, wherein said
2 path selection element defines the at least one path
3 selection rule in a path selection profile.

1 18. The network apparatus of claim 2, wherein said
2 performance enhancing proxy establishes multiple connections
3 of the first type, said performance enhancing proxy
4 including,
5 a compression/encryption element, which compresses
6 and/or encrypts data associated with connections of the
7 first type for transmission across connections of the second
8 type.

1 19. The network apparatus of claim 2, wherein the
2 first connection uses a high layer protocol.

1 20. The network apparatus of claim 2, wherein the
2 first connection uses one of the Transmission Control
3 Protocol (TCP) and the User Datagram Protocol (UDP).

1 21. The network apparatus of claim 2, wherein the
2 second connection is a backbone connection.

1 22. The network apparatus of claim 21, wherein the
2 backbone connection is via a wireless link.

1 23. The network apparatus of claim 22, wherein the
2 wireless link has high latency and high error rate.

1 24. The network apparatus of claim 22, wherein the
2 wireless link is a satellite link.

1 25. The network apparatus of claim 2, wherein said
2 network apparatus is a component of a network gateway.

1 26. The network apparatus of claim 2, wherein said
2 network apparatus is a component of a host.

1 27. The network apparatus of claim 2, wherein said
2 network apparatus is a component of a hub.

1 28. The network apparatus of claim 2, wherein said
2 network apparatus is a component of a VSAT.

1 29. The network apparatus of claim 2, wherein said
2 network apparatus is a component of a router.

1 30. A method, comprising:
2 facilitating communication between a network apparatus
3 and other network entities by performing at least one
4 performance enhancing function.

1 31. The method of claim 30, wherein the network
2 apparatus is connected to other network entities via a first
3 type of connection and other network entities via a second
4 type of connection.

1 32. The method of claim 31, further comprising:
2 establishing multiple connections of the first type
3 associated with different applications; and
4 spoofing some of the multiple connections of the first
5 type based on their associated applications.

1 33. The method of claim 32, wherein said spoofing step
2 only spoofs connections of the first type associated with at
3 least one of applications with high throughput and
4 applications for which reduced startup latency is desired.

1 34. The method of claim 32, wherein said spoofing step
2 assigns spoofing resources, including buffer space and
3 control blocks, to the spoofed connections.

1 35. The method of claim 32, wherein said spoofing step
2 spoofs connections using at least one spoofing rule based on
3 destination address, source address, destination port
4 number, source port number, options, a differentiated
5 services (DS) field or combinations thereof.

1 36. The method of claim 35, wherein said spoofing step
2 defines the at least one spoofing rule in a spoofing
3 profile.

1 37. The method of claim 31, further comprising:
2 establishing multiple connections of the first type;
3 and
4 spoofing acknowledgements (ACKs).

1 38. The method of claim 31, further comprising:
2 establishing multiple connections of the first type;
3 and
4 spoofing a three-way handshake between the network
5 apparatus and another network entity.

1 39. The method of claim 31, further comprising:
2 establishing multiple connections of the first type;
3 and
4 multiplexing multiple connections of the first type
5 onto a single connection of the second type.

1 40. The method of claim 31, further comprising:
2 establishing multiple connections of the first type;
3 and
4 prioritizing connections of the first type to determine
5 what priority level of the connection of the second type,
6 each of the connections of the first type are assigned.

1 41. The method of claim 40, wherein said prioritizing
2 step prioritizes connections using at least one priority
3 rule based on destination address, source address,
4 destination port number, source port number, protocol, a
5 differentiated services (DS) field, type of data contained
6 within the connection or combinations thereof.

1 42. The network apparatus of claim 41, wherein said
2 prioritizing element defines the at least one prioritizing
3 rule in a prioritizing profile.

1 43. The method of claim 31, further comprising:
2 establishing multiple connections of the first type;
3 and
4 selecting a path for data associated with connections
5 of the first type across connections of the second type or
6 connections of other types.

1 44. The method of claim 43, wherein said selection
2 step selects up to N paths ($N > 1$), where the Nth path is
3 selected only if the (N-1)th path fails.

1 45. The method of claim 44, wherein said selection
2 step selects a path using at least one path selection rule
3 based on priority, a destination address, source address,
4 destination port number, source port number, protocol, a
5 differentiated services (DS) field or combinations thereof.

1 46. The method of claim 45, wherein said selection
2 step defines the at least one path selection rule in a path
3 selection profile.

1 47. The method of claim 31, further comprising:
2 establishing multiple connections of the first type;
3 and
4 compressing and/or encrypting data associated with
5 connections of the first type for transmission across
6 connections of the second type.

1 48. The method of claim 31, wherein the first
2 connection uses a high layer protocol.

1 49. The method of claim 31, wherein the first
2 connection uses one of the Transmission Control Protocol
3 (TCP) and the User Datagram Protocol (UDP).

1 50. The method of claim 31, wherein the second
2 connection is a backbone connection.

01 51. The method of claim 50, wherein the backbone
02 connection is via a wireless link.

01 52. The method of claim 51, wherein the wireless link
02 has high latency and high error rate.

01 53. The method of claim 50, wherein the wireless link
02 is a satellite link.

01 54. The method of claim 31, wherein said method is
02 performed in a network gateway.

1 55. The method of claim 31, wherein said method is
2 performed in a host.

1 56. The method of claim 31, wherein said method is
2 performed in a hub.

1 57. The method of claim 31, wherein said method is
2 performed in a VSAT.

1 58. The method of claim 31, wherein said method is
2 performed in a router.

1 59. The method of claim 31, wherein said method is
2 performed in a switch.

1 60. The network apparatus of claim 2, wherein said
2 network apparatus is a component of a switch.

1

ABSTRACT

5

10

15

Method and apparatus for enhancing the performance of a network. The performance enhancing functions described are applicable to a wide variety of communication links, including both slow and fast links, high latency links, and links with low and high error rates. The performance enhancing functions, which may be implemented either singly or in combination, include selective spoofing which allows flexible configuration of which connections should be spoofed, spoofing of the conventional TCP three-way handshake, local data acknowledgement, which allows data windows to increase at local speeds, multiplexing multiple connections across a single connection, data compression/encryption, prioritization, and path selection. The performance enhancing features described are particularly useful for links with high latency and/or high bit error rates.

Figure 1




Figure 2




Figure 4

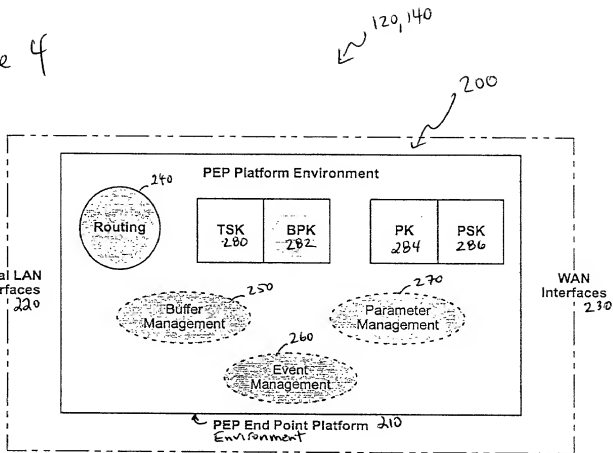
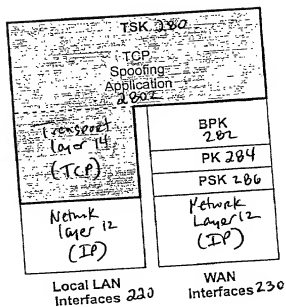



Figure 5



09662072-091400

Figure 6



COMBINED DECLARATION FOR PATENT APPLICATION AND POWER OF ATTORNEY

Page 1 of 2
PD-200053

- X Original
☐ Continuation
☐ Division
☐ Continuation-in-part
☐ Supplemental

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name.

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled: **Time-Offset Distribution to Ensure Constant Satellite Power** the specification of which

(check one) ☒ is attached hereto.
 was filed on _____ as Application Serial No. _____ and (a) [other than supplemental] was amended on or (b) [supplemental] with amendments through _____.

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by an amendment referred to above.

I acknowledge the duty to disclose to the United States Patent and Trademark Office all information known to me to be material to patentability as defined in Title 37, Code of Federal Regulations, §1.56.

I hereby claim foreign priority benefits under Title 35, United States Code, §119(a)-(d) of any foreign application(s) for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate having a filing date before that of the application on which priority is claimed:

Prior Foreign Application(s)

Number	Country	Day/Month/Year Filed	Priority Claimed
_____	_____	_____	<input type="checkbox"/> Yes <input type="checkbox"/> No

I hereby claim foreign priority benefits under Title 35, United States Code, §119(e) of any provisional application(s) for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate having a filing date before that of the application on which priority is claimed:

<u>60/185,553</u> Application Serial No.	<u>February 28, 2000</u> Filing Date	<u>pending</u> Status	Priority Claimed X Yes <input type="checkbox"/> No
<u>60/220,026</u> Application Serial No.	<u>July 21, 2000</u> Filing Date	<u>pending</u> Status	Priority Claimed X Yes <input type="checkbox"/> No

I hereby claim the benefit under Title 35, United States Code, §120 of any United States application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United States Code §112, I acknowledge the duty to disclose to the United States Patent and Trademark Office all information known to me to be material to patentability as defined in Title 37, Code of Federal Regulations, §1.56 which became available between the filing date of the prior application and the national or PCT international filing date of this application:

_____	_____	_____
Application Serial No	Filing Date	Status

Declaration
Page 2 of 2
PD-200053

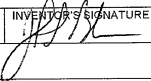
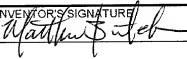
I hereby appoint the following attorneys, or agent and attorneys, to prosecute this application and to transact all business in the Patent and Trademark Office connected therewith:

John T. Whelan
Michael W. Sales
Craig L. Plastrik

Registration No. 32,448
Registration No. 30,213
Registration No. 41,254.

Address all telephone calls to: (301) 428-7172. Address all correspondence to **Customer Number 020991 (Hughes Electronics Corporation, Patent Docket Administration, Bldg. 001, M/S A109, PO Box 956, El Segundo, California 90245-0956).**

I hereby declare that all statement made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such false statements may jeopardize the validity of the application or any patent issued thereon.

FULL NAME OF SOLE OR JOINT INVENTOR John Border	INVENTOR'S SIGNATURE 	DATE 9/14/00
RESIDENCE (CITY AND STATE) Germantown, Maryland		CITIZENSHIP
POST OFFICE ADDRESS 19605 Selby Avenue, Poolesville, Maryland 20837		
FULL NAME OF SOLE OR JOINT INVENTOR Matthew Butehorn	INVENTOR'S SIGNATURE 	DATE 9/12/00
RESIDENCE (CITY AND STATE) Mt. Airy, Maryland		CITIZENSHIP USA
POST OFFICE ADDRESS 1230 Leafy Hollow Circle, Mt. Airy, Maryland 21771		

00852072-01400